

Many-body quantum state tomography with neural networks

Giacomo Torlai,^{1,2} Guglielmo Mazzola,³ Juan Carrasquilla,⁴
Matthias Troyer,^{3,5} Roger Melko,^{1,2} and Giuseppe Carleo^{3,*}

¹*Department of Physics and Astronomy, University of Waterloo, Ontario N2L 3G1, Canada*

²*Perimeter Institute of Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada*

³*Theoretische Physik, ETH Zurich, 8093 Zurich, Switzerland*

⁴*D-Wave Systems Inc., 3033 Beta Avenue, Burnaby BC V5G 4M9, Canada*

⁵*Quantum Architectures and Computation Group,
Station Q, Microsoft Research, Redmond, WA 98052, USA*

The experimental realization of increasingly complex synthetic quantum systems calls for the development of general theoretical methods, to validate and fully exploit quantum resources. Quantum-state tomography (QST) aims at reconstructing the full quantum state from simple measurements, and therefore provides a key tool to obtain reliable analytics. Brute-force approaches to QST, however, demand resources growing exponentially with the number of constituents, making it unfeasible except for small systems. Here we show that machine learning techniques can be efficiently used for QST of highly-entangled states in arbitrary dimension. Remarkably, the resulting approach allows one to reconstruct traditionally challenging many-body quantities – such as the entanglement entropy – from simple, experimentally accessible measurements. This approach can benefit existing and future generations of devices ranging from quantum computers to ultra-cold atom quantum simulators.

Machine-learning (ML) methods have been demonstrated to be particularly powerful at compressing high-dimensional data into low-dimensional representations.^{1,2} Thanks to its intrinsic flexibility, ML is being applied to unravel complex patterns hidden in the most diverse data sources, showing robustness against noise, and receptiveness to generalization. While in the past ML has been mostly applied to data science, it has recently been used to address questions in the physical sciences. Applications to quantum many-body systems have been put forward last year, to classify phases of matter^{3–8}, and to improve the simulation of classical^{9,10} and quantum¹¹ systems.

QST is by itself a data-driven problem,¹² in which we aim to obtain a complete quantum-mechanical description of a system, on the basis of a limited set of experimentally accessible measurements. Key quantum features, such as multi-qubit entanglement, are however challenging to probe directly in current experimental setups.¹³ Finding an efficient method to reliably extract such information from a generic quantum device is therefore important for the development of more powerful quantum simulators.^{14–16} In order to efficiently perform QST, it is necessary to find a compact, and sufficiently general representation of the quantum state to be analyzed. This representation should be general enough to be applied to both low-¹⁷ and high-dimensional systems, and be able to describe highly-entangled quantum states.

In this Letter we show how ML approaches can be used to find such representations. In particular, we argue that suitably-trained artificial neural networks (ANN) offer a natural, efficient, and general way of performing QST driven by experimental data. Our approach is demonstrated on controlled artificial datasets, comprising measurements from several quantum states with a large number of degrees of freedom (qubits, spins, etc...), that are

thus traditionally hard for QST approaches.

The ANN architecture we use in this work is based on restricted Boltzmann machine (RBM) models. RBMs feature a visible layer (describing the physical qubits) and a hidden layer of stochastic binary neurons fully connected with weighted edges to the visible layer. These models have been successfully employed to effectively solve complex many-body problems both for classical¹⁸ and quantum systems¹¹. “Neural quantum state” representations of the many-body wave-function have been shown to be capable of sustaining high entanglement, and to efficiently describe complex topological phases of matter.^{4,19–23} Given these favorable properties, RBM-based quantum states are natural candidates for QST of low and high-dimensional many-body systems.

Our machine-learning approach to QST is realized in two steps. First, an RBM is trained on a dataset consisting of a series of independent density measurements $|\Psi(\mathbf{x}_j)|^2$ in some given basis $\{\mathbf{x}\}$ of the N -body quantum system. If the target state is not positive-definite, additional estimates of the sign (phase) of the wave-function are required. Once trained, the RBM approximates the wave-function underlying the dataset, and can be sampled to determine properties of interest for the physical system under investigation (see Supplementary Material). The accuracy of the reconstruction can be systematically improved by increasing the number of hidden neurons M in the RBM for fixed N , or equivalently the density of hidden units $\alpha = M/N$.¹¹

To demonstrate the power of this approach, we first consider QST of the W state, a paradigmatic N -qubit multipartite entangled wave-function defined as

$$|W\rangle = \frac{1}{\sqrt{N}}(|100\dots\rangle + \dots + |\dots 001\rangle). \quad (1)$$

To mimic the experimental measurements, we generate

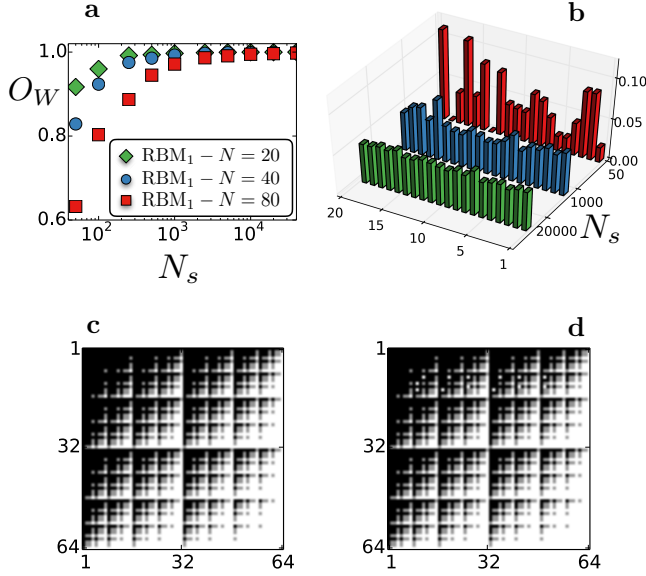


Figure 1. **Tomography of the W state.** **a)** Overlap between the W state and the wave-function generated by the trained $\text{RBM}_{\alpha=1}$ as a function of the number of samples in the training dataset. **b)** Histogram of the occurrence of each of the superposed states in the W state for $N = 20$ qubits. We plot three histograms obtained by sampling a RBM trained on a dataset containing 50 (red), 1000 (blue) and 20000 (green) independent samples. **c-d)** Sign $\mathcal{S}(\mathbf{x})$ of each state $|\mathbf{x}\rangle$ in $|W_X\rangle$ with $N = 12$, where $\mathcal{S} = +1$ (black) if $\langle \mathbf{x} | W_X \rangle > 0$ and $\mathcal{S} = -1$ (white) if $\langle \mathbf{x} | W_X \rangle < 0$. We show the comparison between the exact state (**c**) and the sign learned by the RBM (**d**) (states plotted in grey do not belong to $|W_X\rangle$). RBM tomography allows here to systematically converge to the target W state, when measured in the Z or in the X basis, upon increasing the number of experimental samples.

several datasets with an increasing number of synthetic density measurements obtained by sampling from the W state. These measurements are used to train the RBM model. After the training, we sample states (i.e. configurations of the neurons in the visible layer) from the RBM and build the histogram of the relative frequency of the superimposed states appearing in $|W\rangle$. In Fig. 1(b) we show three histograms obtained with a different number of samples in the training dataset for $N = 20$, and for a fixed density of hidden variables $\alpha = 1$. From the histograms, we see that upon increasing the number of samples each of the N components ($|100\dots\rangle, |010\dots\rangle, \dots$) contribute equally to the wave-function, as expected from the exact W state. To better quantify the quality of our reconstruction we then compute the overlap O_W of the wave-function generated by the RBM with the original W state (see Suppl. Mat.). In Fig. 1(a) we plot the overlap as a function of the number of samples in the training datasets for three different values of N , showing that for a system size substantially larger than what is currently available in experiments,²⁴ an overlap $O_W \sim 1$ can be achieved with a limited number of samples.

Our approach can be extended to the case in which the target state is not positive-definite. For simplicity, we assume that the state can be gauged into a real-valued form,²⁵ and that it is experimentally feasible to estimate the sign of the many-body state on a given set of kets $|\mathbf{x}_1\rangle, \dots, |\mathbf{x}_P\rangle$. For example, this could be realized through a weak measurement in the $\{\mathbf{x}\}$ basis followed by a strong measurement in a chosen complementary basis.²⁶ QST is then realized training two RBM's, one to learn the amplitude $|\Psi(\mathbf{x})|^2$, and an additional one to learn the sign $\mathcal{S}(\mathbf{x})$ of the wave-function (see Suppl. Mat.). After the training, the two networks are combined together to sample the full wave-function $\Psi(\mathbf{x}) = \mathcal{S}(\mathbf{x})|\Psi(\mathbf{x})|$. To demonstrate this approach, we have performed a reconstruction of the W state in the σ_x basis, $|W_X\rangle = H_1 \dots H_N |W\rangle$, where H_i is a Hadamard gate applied to qubit i . The rotated W_X state features a non-trivial sign structure, and is therefore a good test case for QST. We show in Fig. 1(c) the sign $\mathcal{S}(\mathbf{x})$ of each ket $|\mathbf{x}\rangle$ in $|W_X\rangle$ and see that it closely matches the sign learned by the RBM (Fig. 1(d)). The overall scheme has an accuracy similar to the positive-definite case, with the overlap closely approaching unity on system sizes $N \sim 12$ for which a direct, exact sampling from the rotated W_X state is still possible.

We now turn to the case of more complex systems and demonstrate QST for genuine many-body problems. To mimic experimental outcomes, we generate artificial datasets sampling the ground-state of interacting spin models on a lattice. These are directly relevant for quantum simulators based on ultra-cold ions and atoms. In particular we consider the transverse-field Ising model (TFIM) with Hamiltonian

$$\mathcal{H} = - \sum_{\langle ij \rangle} \sigma_i^z \sigma_j^z - h \sum_i \sigma_i^x \quad (2)$$

and the XXZ spin- $\frac{1}{2}$ model, with Hamiltonian

$$\mathcal{H} = \sum_{\langle ij \rangle} \left[\Delta (\sigma_i^x \sigma_j^x + \sigma_i^y \sigma_j^y) + \sigma_i^z \sigma_j^z \right] \quad (3)$$

where the σ_i are Pauli spin operators.

For each Hamiltonian we consider both a 1-dimensional (1d) chain with N sites and a 2-dimensional (2d) square lattice with linear extent L and $N = L^2$. For the TFIM, there exists a quantum critical point h_c in the transverse field dividing a disordered paramagnetic phase ($h > h_c$) from an ordered ferromagnetic phase ($h < h_c$), while the XXZ model undergoes a quantum phase transition at a critical value Δ_c of the anisotropy from a disordered phase ($\Delta > \Delta_c$) to a magnetically ordered phase ($\Delta < \Delta_c$). Synthetic measurements in this case are obtained with standard quantum Monte Carlo (QMC) methods,^{27,28} stochastically sampling the exact ground-state of Hamiltonians in Eqs. (2,3) for different values of the coupling parameters h and Δ . Since in this case the many-body ground-state is positive-definite, our reconstruction scheme does not require extra measurements

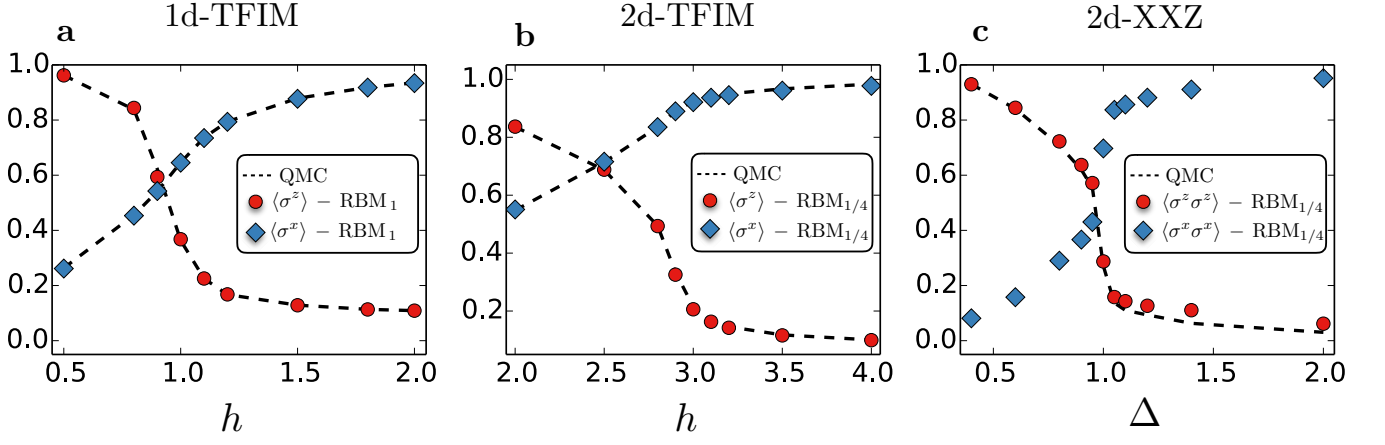


Figure 2. **Reconstruction of few-body observables.** Comparison between different few-body observables computed by sampling the RBM_α (markers) with quantum Monte Carlo simulations (dashed line). Each data point is obtained with a RBM from a network trained on separate datasets. **a)** Diagonal (red circles) and non-diagonal (blue diamonds) magnetization as a function of the transverse field h for the 1d-TFIM with $N = 100$ spins. **b)** Diagonal and non-diagonal magnetization as a function of the transverse field h for the 2d-TFIM on a square lattice with linear size $L = 12$. **c)** Two-point correlation function (diagonal and non-diagonal) between neighboring spins along the diagonal of the square lattice (linear size $L = 12$) for the 2d-XXZ model. RBM tomography allows here to accurately reconstruct both diagonal and off-diagonal observables of the target state.




Figure 3. **Spin-spin correlation function.** Reconstruction of the diagonal spin correlation function $\langle \sigma_i^z \sigma_j^z \rangle$ for the 1d-TFIM with $N = 100$ sites at the critical point $h = 1$. **a)** Direct calculation on the spin configuration in the training dataset, i.e. exact result. **b)** Reconstruction of the correlations by sampling the trained $\text{RBM}_{1/2}$.

of the many-body phase. Once the training is complete, we can test the representational power of the neural networks by computing various observables using the RBMs and comparing with the values obtained through QMC simulations.¹⁸

We begin by considering few-body observables, such as magnetization and spin correlations. For the TFIM we first look at the longitudinal σ^z magnetization. Because this observable is diagonal in the basis adopted for training, the problem reduces to a purely classical one¹⁸ and we can easily compute the expectation value from the configurations of the visible layer, sampled from the trained RBM. As shown in Fig. 2 (a-b) for one and two dimensions, the RBMs can reproduce the average val-

ues with high accuracy. We then consider the case of the transverse σ^x magnetization, which is off-diagonal in the training basis. One can obtain the expectation value by computing the local estimate of the observable as in variational Monte Carlo simulations (see Suppl.). Even though reproducing this off-diagonal observable is non-trivial, the RBM can nonetheless generate values in agreement with QMC calculations. For the XXZ model, we show in Fig. 2 (c) the expectation values of the diagonal $\sigma_a^z \sigma_b^z$ and off-diagonal $\sigma_a^x \sigma_b^x$ spin correlations, with a and b being neighbors along the lattice diagonal. Finally, we consider the full spin-spin $\sigma_i^z \sigma_j^z$ correlation function for the 1d-TFIM, which involves non-local correlations. In Fig. 3 we show the reconstruction of the correlation




Figure 4. **Entanglement entropy.** Comparison of the second Renyi entropy as a function of the subsystem size ℓ between the RBMs (markers) and results obtained via exact diagonalization (dashed lines). **a)** 1d-TFIM at three different values of the transverse field, $h = 0.8$ (red squares), $h = 1.0$ (blue circles) and $h = 1.2$ (green diamonds) with $N = 20$ spins. **b)** 1d Heisenberg model ($\Delta = 1.0$) with $N = 20$ spins (blue circles).

function using the RBM **(a)** closely matching the exact result obtained by direct computation from the spin states making up the training dataset **(b)**.

To further assess the capabilities of our approach, we finally turn to the entanglement entropy, which is a highly non-local quantity particularly challenging for direct experimental observations. It provides important information on the universal behavior of interacting many-body systems and it is of central interest in condensed matter physics and quantum information theory. The entanglement entropy has been computed numerically for various physical systems,²⁹ however its highly non-local nature makes it very challenging to measure in experiments.¹³ Following the method proposed here, we can obtain an estimate of this quantity given only simple measurements of the density, which are more accessible with current ex-

perimental advances³⁰. Given a bipartition of the physical system, we consider in particular the second Renyi entropy defined as $S_2(\rho_A) = -\log(\text{Tr}(\rho_A^2))$, with the subsystem ρ_A of varying size. We estimate S_2 by employing an improved ratio trick sampling²⁹ using the wavefunction generated by the RBM. In Fig. 4 we show the entanglement entropy for the 1d-TFIM **(a)** for three values of the transverse field, and for the critical ($\Delta = 1$) 1d-XXZ model **(b)**. In both instances we took a chain with $N = 20$ spins and plot the entanglement entropy as a function of the subsystem size $\ell \in [1, N/2]$. The values obtained with the RBM (markers) are compared with results from exact diagonalization (dashed lines), with an overall good agreement.

To conclude, we have demonstrated that ML tools can be efficiently used to reconstruct complex many-body quantum states from a limited number of experimental measurements. Our scheme is general enough to be efficiently applied to a variety of quantum devices for which current approaches demand exponentially large resources. These include QST of highly-entangled quantum circuits, adiabatic quantum simulators,³¹ experiments with ultra-cold atoms and ions traps in higher dimensions. Our approach can be used to directly validate quantum computers and simulators, as well as to indirectly reconstruct quantities which are experimentally challenging for a direct observation. For example, we anticipate that the current generation of quantum microscopes could substantially benefit from neural-quantum states QST. In particular, we predict that the use of our approach for bosonic ultra-cold atoms experiments would allow for the determination of the entanglement entropy on systems substantially larger than those currently accessible with quantum interference techniques.¹³

Acknowledgements

We thank H. Carteret and B. Kulchytskyy for useful discussions. GT and RGM acknowledge support from NSERC, the Canada Research Chair program, the Ontario Trillium Foundation, and the Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported through Industry Canada and by the Province of Ontario through the Ministry of Research & Innovation. GC, GM and MT acknowledge support from the European Research Council through ERC Advanced Grant SIMCOFE, and the Swiss National Science Foundation through NCCR QSIT. Simulations were performed on resources provided by SHARCNET, and by the Swiss National Supercomputing Centre CSCS.

* gcarleo@phys.ethz.ch

¹ Hinton, G. E. & Salakhutdinov, R. R. Reducing the Dimensionality of Data with Neural Networks. *Science* **313**,

- 504–507 (2006). URL <http://science.sciencemag.org/content/313/5786/504>.
- ² LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015). URL <http://www.nature.com/nature/journal/v521/n7553/full/nature14539.html>.
 - ³ Wang, L. Discovering phase transitions with unsupervised learning. *Physical Review B* **94**, 195105 (2016). URL <http://link.aps.org/doi/10.1103/PhysRevB.94.195105>.
 - ⁴ Carrasquilla, J. & Melko, R. G. Machine learning phases of matter. *Nature Physics* **advance online publication** (2017). URL <http://www.nature.com/nphys/journal/vaop/ncurrent/full/nphys4035.html>.
 - ⁵ Broecker, P., Carrasquilla, J., Melko, R. G. & Trebst, S. Machine learning quantum phases of matter beyond the fermion sign problem. *arXiv:1608.07848* (2016). URL <http://arxiv.org/abs/1608.07848>.
 - ⁶ Ch’ng, K., Carrasquilla, J., Melko, R. G. & Khatami, E. Machine Learning Phases of Strongly Correlated Fermions. *arXiv:1609.02552* (2016). URL <http://arxiv.org/abs/1609.02552>.
 - ⁷ van Nieuwenburg, E. P. L., Liu, Y.-H. & Huber, S. D. Learning phase transitions by confusion. *Nature Physics* **advance online publication** (2017). URL <http://www.nature.com/nphys/journal/vaop/ncurrent/full/nphys4037.html>.
 - ⁸ Ohtsuki, T. & Ohtsuki, T. Deep learning the quantum phase transitions in random electron systems: Applications to Three Dimensions. *arXiv:1612.04909* (2016). URL <http://arxiv.org/abs/1612.04909>.
 - ⁹ Huang, L. & Wang, L. Accelerated Monte Carlo simulations with restricted Boltzmann machines. *Physical Review B* **95**, 035105 (2017). URL <http://link.aps.org/doi/10.1103/PhysRevB.95.035105>.
 - ¹⁰ Liu, J., Qi, Y., Meng, Z. Y. & Fu, L. Self-learning Monte Carlo method. *Physical Review B* **95**, 041101 (2017). URL <http://link.aps.org/doi/10.1103/PhysRevB.95.041101>.
 - ¹¹ Carleo, G. & Troyer, M. Solving the quantum many-body problem with artificial neural networks. *Science* **355**, 602–606 (2017). URL <http://science.sciencemag.org/content/355/6325/602>.
 - ¹² Vogel, K. & Risken, H. Determination of quasiprobability distributions in terms of probability distributions for the rotated quadrature phase. *Physical Review A* **40**, 2847 (1989). URL <http://journals.aps.org/pra/abstract/10.1103/PhysRevA.40.2847>.
 - ¹³ Islam, R. *et al.* Measuring entanglement entropy in a quantum many-body system. *Nature* **528**, 77–83 (2015). URL <http://www.nature.com/nature/journal/v528/n7580/full/nature15750.html>.
 - ¹⁴ Bloch, I., Dalibard, J. & Zwerger, W. Many-body physics with ultracold gases. *Reviews of Modern Physics* **80**, 885–964 (2008).
 - ¹⁵ Blatt, R. & Roos, C. F. Quantum simulations with trapped ions. *Nature Physics* **8**, 277–284 (2012). URL <http://www.nature.com/nphys/journal/v8/n4/abs/nphys2252.html>.
 - ¹⁶ Shulman, M. D. *et al.* Demonstration of Entanglement of Electrostatically Coupled Singlet-Triplet Qubits. *Science* **336**, 202–205 (2012). URL <http://science.sciencemag.org/content/336/6078/202>.
 - ¹⁷ Cramer, M. *et al.* Efficient quantum state tomography. *Nature communications* **1**, 149–149 (2009). URL <http://www.nature.com/articles/ncomms1147>.
 - ¹⁸ Torlai, G. & Melko, R. G. Learning thermodynamics with Boltzmann machines. *Physical Review B* **94**, 165134 (2016). URL <http://link.aps.org/doi/10.1103/PhysRevB.94.165134>.
 - ¹⁹ Deng, D.-L., Li, X. & Sarma, S. D. Exact Machine Learning Topological States. *arXiv:1609.09060* (2016). URL <http://arxiv.org/abs/1609.09060>.
 - ²⁰ Deng, D.-L., Li, X. & Sarma, S. D. Quantum Entanglement in Neural Network States. *arXiv:1701.04844* (2017). URL <http://arxiv.org/abs/1701.04844>.
 - ²¹ Gao, X. & Duan, L.-M. Efficient Representation of Quantum Many-body States with Deep Neural Networks. *arXiv:1701.05039* (2017). URL <http://arxiv.org/abs/1701.05039>.
 - ²² Chen, J., Cheng, S., Xie, H., Wang, L. & Xiang, T. On the Equivalence of Restricted Boltzmann Machines and Tensor Network States. *arXiv:1701.04831* (2017). URL <http://arxiv.org/abs/1701.04831>.
 - ²³ Huang, Y. & Moore, J. E. Neural network representation of tensor network and chiral states. *arXiv:1701.06246* (2017). URL <http://arxiv.org/abs/1701.06246>. ArXiv: 1701.06246.
 - ²⁴ Wang, X.-L. *et al.* Experimental Ten-Photon Entanglement. *Physical Review Letters* **117**, 210502 (2016). URL <http://link.aps.org/doi/10.1103/PhysRevLett.117.210502>.
 - ²⁵ The RBM tomography can be however extended also to more general complex-valued states.
 - ²⁶ Lundeen, J. S., Sutherland, B., Patel, A., Stewart, C. & Bamber, C. Direct measurement of the quantum wavefunction. *Nature* **474**, 188–191 (2011). URL <http://www.nature.com/nature/journal/v474/n7350/full/nature10120.html>.
 - ²⁷ Rieger, H. & Kawashima, N. Application of a continuous time cluster algorithm to the two-dimensional random quantum ising ferromagnet. *The European Physical Journal B - Condensed Matter and Complex Systems* **9**, 233–236 (1999). URL <http://dx.doi.org/10.1007/s100510050761>.
 - ²⁸ Evertz, H. G. The loop algorithm. *Advances in Physics* **52**, 1–66 (2003). URL <http://dx.doi.org/10.1080/0001873021000049195>.
 - ²⁹ Hastings, M. B., González, I., Kallin, A. B. & Melko, R. G. Measuring Renyi Entanglement Entropy in Quantum Monte Carlo Simulations. *Physical Review Letters* **104**, 157201 (2010). URL <http://link.aps.org/doi/10.1103/PhysRevLett.104.157201>.
 - ³⁰ Bakr, W. S. *et al.* Probing the superfluid-to-mott Insulator Transition at the Single-Atom Level. *Science* **329**, 547–550 (2010). URL <http://science.sciencemag.org/content/329/5991/547>.
 - ³¹ Johnson, M. W. *et al.* Quantum annealing with manufactured spins. *Nature* **473**, 194–198 (2011). URL <http://www.nature.com/nature/journal/v473/n7346/full/nature10012.html>.

SUPPLEMENTARY MATERIALS

Appendix A: Training stage

A restricted Boltzmann machine (RBM) is a probabilistic model defined on a graph with two layers of

stochastic binary neurons, a visible layer \mathbf{x} and a hidden layer \mathbf{h} . The RBM associates to this graph structure a probability distribution given by the Boltzmann distribution

$$p_{\lambda}(\mathbf{x}, \mathbf{h}) = \frac{1}{Z_{\lambda}} \exp \left[\sum_j b_j x_j + \sum_i c_i h_i + \sum_{ij} W_{ij} h_i x_j \right], \quad (\text{A1})$$

where Z_{λ} is the partition function and the network parameters λ are the weights \mathbf{W} connecting the two layers and the biases \mathbf{b} and \mathbf{c} coupled to N visible and M hidden neurons, respectively. The distribution (of interest) over the visible layer is obtained by marginalization over the hidden degrees of freedom

$$p_{\lambda}(\mathbf{x}) = \frac{1}{Z_{\lambda}} \exp \left[\sum_j b_j x_j + \sum_i \log(1 + e^{c_i + \sum_j W_{ij} x_j}) \right]. \quad (\text{A2})$$

The goal of the training is to obtain a set of weights generating a distribution as close as possible to the distribution underlying the physical system, i.e. $|\Psi(\mathbf{x})|^2$. To this end, we first build a dataset $\mathcal{D} = \{\mathbf{x}^{(\ell)}\}$ by projecting the wave-function on the basis $\{\mathbf{x}\}$ and then tune the network parameters λ to minimize the Kullback-Leibler (KL) divergence between the two distributions, defined as

$$\text{KL}(|\Psi(\mathbf{x})|^2 || p_{\lambda}(\mathbf{x})) \equiv \sum_{\mathbf{x}} |\Psi(\mathbf{x})|^2 \log \frac{|\Psi(\mathbf{x})|^2}{p_{\lambda}(\mathbf{x})} \quad (\text{A3})$$

$$\propto - \sum_{\mathbf{x} \in \mathcal{D}} \log p_{\lambda}(\mathbf{x}), \quad (\text{A4})$$

where we have omitted a constant entropy term on the right hand side. To solve the optimization problem we employ simple gradient descent algorithm

$$\Delta \lambda_j = -\eta \nabla_{\lambda_j} \text{KL}(|\Psi(\mathbf{x})|^2 || p_{\lambda}(\mathbf{x})), \quad (\text{A5})$$

where η is called the learning rate. We point out that, because of the partition function Z_{λ} is not directly accessible, the gradient of the KL divergence must be evaluated using a Markov-chain Monte Carlo simulation.

We now turn to the case of a non positive-definite (but real) target wave-function. In such case, we can write $\Psi(\mathbf{x}) = \mathcal{S}(\mathbf{x})|\Psi(\mathbf{x})|$ where $\mathcal{S}(\mathbf{x}) = \pm 1$. As before, a RBM can be trained to learn the distribution $|\Psi(\mathbf{x})|^2$. Moreover, an additional RBM must be trained to learn the sign $\mathcal{S}(\mathbf{x})$. To achieve this task, we slightly modify the network structure by adding a new “sign” layer \mathbf{y} where we encode the sign using 2 bits (i.e. neurons) and couple the new layer to the hidden layer with a weight

matrix \mathbf{U} and fields \mathbf{d} . After integrating out the hidden layer from the probability distribution over the whole graph, the marginal distribution over the visible and sign layers now reads

$$p_{\mu}(\mathbf{x}, \mathbf{y}) = \frac{1}{Z_{\mu}} \exp \left[\sum_j b_j x_j + \sum_k d_k y_k + \sum_i \log \left(1 + e^{c_i + \sum_j W_{ij} x_j + \sum_k U_{ik} y_k} \right) \right], \quad (\text{A6})$$

where μ is a different set of parameters. Since for the sign we are not interested in sampling the machine, but only in computing the sign $\mathcal{S}(\mathbf{x})$ for a given input state \mathbf{x} , instead of the KL divergence we train the network by maximizing the log-likelihood $\mathcal{L} = \sum_{\mathbf{x}, \mathbf{y} \in \mathcal{D}} \log p_{\mu}(\mathbf{y} | \mathbf{x})$ over a dataset $\mathcal{D} = \{\mathbf{x}^{(\ell)}, \mathbf{y}^{(\ell)}\}$. We point out that the conditional probability

$$p_{\mu}(y_q | \mathbf{x}) = \frac{e^{d_q + \sum_i \log(1 + e^{c_i + \sum_j W_{ij} x_j + U_{iq}})}}{\sum_k e^{d_k + \sum_i \log(1 + e^{c_i + \sum_j W_{ij} x_j + U_{ik}})}} \quad (\text{A7})$$

does not contain the partition function and thus can be computed exactly.

Appendix B: Measurement stage

Once the training is complete, we obtain an approximate representation of the system wave-function $\Psi(\mathbf{x})$ given by $\Psi_{\lambda\mu}(\mathbf{x}) = \mathcal{S}_{\mu}(\mathbf{x}) \sqrt{\tilde{p}_{\lambda}(\mathbf{x})}$ (where $\tilde{p}_{\lambda}(\mathbf{x}) = Z_{\lambda} p_{\lambda}(\mathbf{x})$ is the un-normalized distribution). If $\Psi(\mathbf{x})$ is positive-definite, we only train the RBM to learn the amplitude and thus $\mathcal{S}_{\mu}(\mathbf{x}) = 1$. On the other hand, for the case of non positive-definite wave-function, we have

$$\mathcal{S}_{\mu}(\mathbf{x}) = \Theta(0.5 - p_{\mu}(\mathbf{y}_1 | \mathbf{x})) - \Theta(0.5 - p_{\mu}(\mathbf{y}_0 | \mathbf{x})), \quad (\text{B1})$$

where $\Theta(\odot)$ is the Heaviside theta. The full wave-function $\Psi_{\lambda\mu}(\mathbf{x})$ can then be used to compute all the observables of interests. For the case of the W_Z and W_X states, we computed the overlap $O = \frac{\langle W_{Z/X} | \Psi_{\lambda\mu} \rangle}{\langle \Psi_{\lambda\mu} | \Psi_{\lambda\mu} \rangle}$. For the quantum spins model, the expectation value of any operator $\mathcal{O}^{(D)}$ diagonal on the basis $\{\mathbf{x}\}$ can be computed as

$$\langle \mathcal{O} \rangle_{\tilde{p}} = \frac{1}{\sum_{\mathbf{x}} \tilde{p}_{\lambda}(\mathbf{x})} \sum_{\mathbf{x}} \tilde{p}_{\lambda}(\mathbf{x}) \mathcal{O}_{\mathbf{x}\mathbf{x}}, \quad (\text{B2})$$

On the other hand, for off-diagonal observables $\mathcal{O}^{(\text{ND})}$, the expectation value is obtained by $\langle \mathcal{O}^{(\text{ND})} \rangle = \langle \mathcal{O}_L^{(\text{ND})} \rangle_{\tilde{p}}$, where the so-called local estimate of \mathcal{O} is given by

$$\mathcal{O}_L^{(\text{ND})}(\mathbf{x}) = \sum_{\mathbf{x}'} \sqrt{\frac{\tilde{p}_{\lambda}(\mathbf{x}')}{\tilde{p}_{\lambda}(\mathbf{x})}} \mathcal{O}_{\mathbf{x}\mathbf{x}'}^{(\text{ND})}. \quad (\text{B3})$$